

WEST Search History

DATE: Friday, March 19, 2004

Hide?	<u>Set</u> <u>Name</u>	<u>Query</u>	<u>Hit</u> <u>Count</u>
		<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ</i>	
<input type="checkbox"/>	L24	merkey.in.	11
<input type="checkbox"/>	L23	l16 and network\$	82
		<i>DB=PGPB,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ</i>	
<input type="checkbox"/>	L22	l16 and network\$	23
<input type="checkbox"/>	L21	L20	0
		<i>DB=USPT; PLUR=YES; OP=ADJ</i>	
<input type="checkbox"/>	L20	L19	50
		<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ</i>	
<input type="checkbox"/>	L19	l16 and l6	64
		<i>DB=PGPB,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ</i>	
<input type="checkbox"/>	L18	L16	65
		<i>DB=USPT; PLUR=YES; OP=ADJ</i>	
<input type="checkbox"/>	L17	L16	115
		<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ</i>	
<input type="checkbox"/>	L16	L15 and l7	180
<input type="checkbox"/>	L15	l1 and l11 and l4	3145
<input type="checkbox"/>	L14	L13 and l11	3
<input type="checkbox"/>	L13	l3 and l7	251
<input type="checkbox"/>	L12	L11 and l5	22
<input type="checkbox"/>	L11	((raid or raids or redundan\$ NEAR2 array\$) or (disk or disc or disks or discs)).ab,ti.	717467
<input type="checkbox"/>	L10	l8 and L9	31
<input type="checkbox"/>	L9	l5 and l6	76
<input type="checkbox"/>	L8	l5 and (axis or ordinal)	111
<input type="checkbox"/>	L7	axis	2596291
<input type="checkbox"/>	L6	tape or tapes	635419
<input type="checkbox"/>	L5	l3 and L4	270
<input type="checkbox"/>	L4	(disk or disc or disks or discs)	1570604
<input type="checkbox"/>	L3	l1 and L2	714
<input type="checkbox"/>	L2	scalable coherent interfac\$ or sci	291085
<input type="checkbox"/>	L1	raid or raids or redundan\$ NEAR2 array\$	23576

END OF SEARCH HISTORY

h e b b cg b chh e e f c e b e

Print Selection

Section: ☒ Page(s): Print Copy:

Select?	Document ID	Section(s)	Page(s)	# Pages to print	Database
<input checked="" type="checkbox"/>	6601138	all	all	17	PGPB,USPT,USOC,EPAB,JPAB,DWPI
<input checked="" type="checkbox"/>	5862313	all	all	15	PGPB,USPT,USOC,EPAB,JPAB,DWPI
<input checked="" type="checkbox"/>	6338112	all	all	22	PGPB,USPT,USOC,EPAB,JPAB,DWPI
<input checked="" type="checkbox"/>	5805788	all	all	12	PGPB,USPT,USOC,EPAB,JPAB,DWPI

Building Room Printer

WEST Search History

DATE: Friday, March 19, 2004

Hide?	<u>Set</u> <u>Name</u>	<u>Query</u>	<u>Hit</u> <u>Count</u>
		<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ</i>	
<input type="checkbox"/>	L12	L11 and l5	22
<input type="checkbox"/>	L11	((raid or raids or redundan\$ NEAR2 array\$) or (disk or disc or disks or discs)).ab,ti.	717467
<input type="checkbox"/>	L10	l8 and L9	31
<input type="checkbox"/>	L9	l5 and l6	76
<input type="checkbox"/>	L8	l5 and (axis or ordinal)	111
<input type="checkbox"/>	L7	axis	2596291
<input type="checkbox"/>	L6	tape or tapes	635419
<input type="checkbox"/>	L5	l3 and L4	270
<input type="checkbox"/>	L4	(disk or disc or disks or discs)	1570604
<input type="checkbox"/>	L3	l1 and L2	714
<input type="checkbox"/>	L2	scalable coherent interfac\$ or sci	291085
<input type="checkbox"/>	L1	raid or raids or redundan\$ NEAR2 array\$	23576

END OF SEARCH HISTORY

[First Hit](#) [Fwd Refs](#)☐ [Generate Collection](#) [Print](#)

L9: Entry 64 of 76

File: USPT

Dec 7, 1999

DOCUMENT-IDENTIFIER: US 5999930 A

TITLE: Method and apparatus for distributed control of a shared storage volume

Brief Summary Text (10):

It is common in computer systems for calculation intensive processes to take an enormous amount of time to come to completion. In multi-processor systems, the calculation may be expedited by the distribution of the process over the processors present in the computer. This is possible-because the processors that are resident in a single machine share the same physical resources (e.g. PCI, Memory, Disk, and CPU busses) that the processes are based on, in a direct attached manner. The disadvantage to this approach is that it requires very expensive and uncommon hardware equipment, and software has to be written to take advantage of this capability on a particular hardware model. In addition, the processing power is strictly limited to the number of CPU's that can be put in a single computer system. The single machine, multi-processor approach is further limited by the severe bottlenecks that can be created as the physical resources are divided among an increasing number of CPU's.

Detailed Description Text (2):

FIG. 1A is a hardware block diagram of a prior art client server network. The network comprises clients 54-56 [labeled respectively as clients A-B], server 60 and storage volumes 64 and 66. Storage volume 66 is a compact disk. Storage volume 64 is a redundant array of disks (RAID). Each storage volume has a file directory associated with it. Storage volume 64 contains a file directory 62. Each of clients 54-56 contains respectively cache copies 50-52 of file directory 62.

Detailed Description Text (5):

Shown in software module block 82, are the software modules which are enabled during power on. The module associated with the OS are the access control module 96, the command receipt module 98, the scheduling module 102 and the logical to physical translation and space allocation module 110. The modules are the screen driver 104, the CD driver 106 and the RAID storage driver 108. Finally, application module 100 is enabled by the start-up file 94.

Detailed Description Text (6):

FIG. 1C is a block diagram showing the functional input output I/O relationship between the software modules shown in FIGS. 1B-C includes a client application 150 communicating with server 60, which in turn communicates with either of storage devices 64 or 66, as shown initially in FIG. 1A. The file server 60 comprises a logical I/O block 158 and physical I/O block 160. Logical I/O block 158 comprises command receipt module 98, access control modules 96 and file directories 166. Physical I/O block 160 comprises scheduling module 102, logical to physical translation and space allocation module 110 and drivers 104-108 associated with respectively the monitor, the CD and the RAID storage device. RAID storage device 64 includes a file directory 62 which contains the directory tree and access rights for all data located on storage device 64.

Detailed Description Text (7):

In operation, a client application 150 communicates via server 60 with either of storage devices 64 and 66. Commands received from client application 150 includes:

create, delete, open, close, read, and write. Each of these commands is accompanied by a logical file name. A command is passed from the command receipt module 98 to the access control module 96. The access control modules 96 each interface with a corresponding one of file directories 166. There is generally one file directory for each physical volume. Additionally a volume could exist accross multiple physical storage devices as in software striping. A physical volume can be a single storage device or a partition on a storage device. A file directory will be discussed in greater detail in connection with FIG. 4, and generally includes a directory tree of logical file names, and access privileges associated with each file name. The access control module 96 determines on the basis of access privileges contained in file directories 166 whether or not a specific command should be implemented. For example, if an application command comprises a write request to a file or volume, which is listed as read only, then that request will not be implemented by the access control module. Alternately, when a command is received that the access control module will implement then the data associated with that command is passed directly to the scheduling module 102. The physical I/O block is the primary interface with the environment outside the computer system. The physical I/O deals with blocks of data that are exchanged with disk or tape systems. Thus, it is concerned with the placement of those blocks on the secondary storage device and with the buffering of those blocks in main memory, the scheduling module is concerned with scheduling disk or tape accesses to optimize performance. I/O buffers are assigned and secondary memory is allocated at this level. From the scheduling module 102, a command to, e.g., read data from a specific logical location is passed to logical to physical translation and space allocation module 110. This module is responsible for translating a logical location into a actual physical location on a specific storage device. A file for example, may not be stored on a contiguous portion of a storage device. Instead, it may be stored on fragmented segments. This module then performs the translation of a logical file name to a series of physical segments. This module is also responsible for space allocation. Space allocation in the sense that when a write commands is received the actual physical locations to which the data is written, is determined by module 110. The final software interface between the server 60 and the storage device is via an appropriate device driver. Each device driver handles one device type. The device driver is the only part of the operating system that knows how many registers the disk controller has and what they are used for. It alone knows about sectors, tracks, cylinders, head-arm motion, interleave factors, motor drives, head settling times and all other mechanics of making the storage device work properly. In general terms, the job of a device driver is to accept abstract requests from the device-independent software above it and see to it that the request is executed.

Detailed Description Text (8):

FIG. 2A is a hardware block diagram of a preferred embodiment of the current invention in which clients may be directly attached to one or more shared storage volumes, without loss of coherency of the data contained on those volumes. Clients 54-56 storage devices 64-66 and server 60 are shown. RAID storage device 64 includes a file directory 204, an access control table 206 and a volume control table 208. Clients 54-56 each implement processes for distributed management of storage volumes 64-66. Client 54 implements process 214 and client 56 implements process 216. Clients 54-56 include resident copies 200-202 of file directory 62.

Detailed Description Text (9):

Server 60 as indicated by x marker 212 is no longer needed in order to service a request from either client 54 or 56 to the storage volumes 64-66. Instead, clients 54-56 are directly connected to storage volumes 64-66. In a preferred embodiment the connection is based on either the Fibre Channel ANSI Standard X3.230 and/or the SCSI-3 ANSI Standard X3.270, although other media connection systems may be supported. Other embodiments of the present invention may be adapted to support other methods of connecting the storage 802 and hosts 804, such as embodiments utilizing Fast-40 (Utra-SCSI), Serial Storage Architecture (SSA), "FireWire" IEEE

Standard P1394, Asynchronous Transfer Mode (ATM), or Scalable Coherent Interface (SCI) IEEE Standard 1596-1992, or, some combination of the above, among other possibilities. The fibre channel architecture provides high speed interface links to both serial communications and storage I/O. These links can transmit data using industry standard interfaces at 10 to 100 times the speed of currently-existing architectures. Fibre channel is part of a very recent revolution in data communications caused by incorporating the serial transmission line, design techniques and technologies into applications and computer architecture that have traditionally used parallel bus-base types of data transport. Transmission lines have superior characteristics for data transport and bus-based designs, allowing signals to be transmitted error-free over greater distances and/or at higher data rates. Fibre channel supports a flexible topology including point-to-point links, packet-switching, and shared-media loop topologies. In a preferred embodiment, a fibre channel network links clients 54-56 to both of storage volumes 64-66. Storage volume 64 is equipped with striping software which reads and writes data concurrently across multiple disks. This striping capability allows the bandwidth of the network medium, i.e., fibre channel to be equalled by the bandwidth of the storage device. This is in spite of the fact that individual disks on the storage device can only provide data at a fraction of the bandwidth of the network connection. By striping a file across multiple disks the bandwidth constraints imposed by the physical limits on the head-arm are overcome by allowing multiple disks to operate in parallel to provide data drives.

Detailed Description Text (12):

At power on CPU 250 uploads the OS kernel from boot ROM 254 and reads the remaining portion of the operating system 260 from the local storage device 258, e.g., a disk drive. During the latter stages of the operating system enablement, device drivers 262 and start-up file 264 are enabled. Software module block 252 shows the software modules that are enabled during the power on. During the enablement of the operating system 260, the command receipt module 98, the access control modules 96, the scheduling module 102, the logical to physical translation and space allocation module 110, and the converter modules 268 are enabled. During the enablement of the device drivers 262, the screen driver 104, the CD driver 106 and the RAID driver 270 are enabled. During the initiation of the start-up file, the volume control application 266 and a resident application 100 are enabled.

Detailed Description Text (17):

FIG. 3B shows the data structures associated with the volume control tables 208. Specifically, tables 208A and 208B are shown. There is one volume control table for each volume, which supports distributed control by multiple writers. The fields within volume control record 208A are all identified with the suffix "A." The fields identified with volume control table 208B all contain a "B" suffix. Field 390 identifies the logical volume associated with a specific volume control table. Field 390A indicates that RAID storage device 64 [see FIG. 2A] is associated with volume control table 208A. Field 390B indicates that the CD-ROM 66 [see FIG. 2A] is associated with volume control table 208B. Field 392 contains a boolean TRUE/FALSE indicating whether or not a specific volume is write enabled. Field 392A indicates that the RAID device 64 [see FIG. 2A] is write enabled. Field 392B indicates that CD-ROM 66 [see FIG. 2A] is not write enabled. Field 394 indicates which client currently has write access to a specific volume. Field 394-A indicates that client 54 [see FIG. 2A] currently has write access to RAID storage device 64. Field 394B indicates that no client has write access to CD-ROM 66 [see FIG. 2A]. Field 388 indicate which clients have mount access privileges for each specific volume. A boolean true indicates that the client can mount the volume. A boolean false indicates the opposite. Field 396 indicates for each client as to the ability to request a change to its current volume settings. A boolean false indicates a client is not locked out from making change requests such as read-only to read-write (or vice versa). A boolean true indicates a client is locked out from making change requests. Field 384 is a boolean true/false indicating whether or not a client with read only privileges will be updated when changes are made to the volume by other

clients with respect to a specific volume. Fields 386 are time stamps indicating the last time at which a client received an updated copy of a file directory 62 [See FIG. 2A]. Field 382 is a time stamp indicating the last modification time for a specific volume by any client. By comparing the last modification time field 386 to the volume modification time field 382 the processes 214-216 [see FIG. 2A] can determine when a client with auto update privileges is in need of a file directory refresh.

Detailed Description Text (178):

FIG. 10 shows the software modules on each client for managing the distributed database. Client 54 and RAID storage device 64 are shown. RAID storage device 64 contains a database and locking fields 804. The client 54 includes run time supervisor 900, transaction routines 902, command and path intercept 904, buffer manager 908, lock manager 906 and lock rerouter 910.